Aegis Orbital Compute Node (AOCN)

LEO Compute Architecture Constrained by Power Generation, Heat Rejection, and Serviceability

Program Positioning

The **Aegis Orbital Compute Node (AOCN)** is a modular compute platform intended for deployment in **Low Earth Orbit (LEO)**. The architecture is defined by first-order physical constraints: electrical power generation, waste heat rejection, radiation environment, and on-orbit servicing.

AOCN is not intended to replicate terrestrial data centers in orbit. It is structured as an orbital compute accelerator for workloads where power density, thermal isolation, or space adjacency dominate over latency, interactive access, or geographic proximity.

The concept is grounded in current launch capabilities, existing space-qualified power and thermal technologies, and commercially available compute hardware. No speculative propulsion systems, materials, or unproven physics are assumed.

Design Constraints and Assumptions

AOCN is organized around the following constraints:

1. **Thermal scaling dominates**
   Sustained compute output is bounded by radiator area, operating temperature, and heat transport capacity. Compute density does not set system scale.

2. **Infrastructure separation**
   Power generation, thermal control, shielding, and communications are treated as persistent infrastructure. Compute hardware is assumed to have a shorter refresh cycle and is designed to be replaceable.

3. **Incremental deployment**
   Nodes are designed to be deployed, serviced, and upgraded independently, avoiding monolithic structures with large fault domains and high replacement risk.

Baseline Orbit and Operating Regime

- **Orbit:** 500–600 km circular LEO

- **Inclination:** 51.6° (ISS-class) or sun-synchronous variants
- **Radiation Environment:** Magnetosphere-protected; manageable single-event upset (SEU) rates
- **Design Life:** 10–15 years per node with modular refresh

LEO is selected to balance radiation exposure, servicing accessibility, communications latency, and launch cost within current operational norms.

System Architecture Overview

1. Structural Core

- Cylindrical or hex-prismatic pressure-tolerant module
- Aluminum-lithium primary structure
- Internal racks mounted on vibration-isolated frames
- External attachment points for radiators, solar wings, and communications booms

No rotating elements are required. Artificial gravity is unnecessary for electronic systems and would complicate thermal and structural design.

2. Compute Payload

- **Hardware class:** Commercial GPUs, accelerators, and CPUs (COTS)
- **Form factor:** Ruggedized rack units with conformal coating
- **Fault tolerance:**
    - ECC memory throughout
    - Continuous memory scrubbing
    - Checkpointed workloads
    - Graceful degradation at the node level

The system is designed under the assumption that single-event upsets occur and are handled at the software and system level. This approach mirrors long-duration ISS avionics practices applied to modern compute payloads.

3. Thermal Management and Heat Rejection (Critical Path)

All electrical input power ultimately becomes waste heat.

The thermal architecture consists of:

- Cold plates at the compute rack level
- Pumped liquid cooling loops
- Heat exchangers feeding deployable radiator panels

At steady state, essentially all electrical input power is rejected as heat; thermal subsystem performance therefore sets the upper bound on continuous operation.

Radiator operating regime:

- Operating temperature: ~350–450 K
- Heat rejection: ~500–1,000 W/m² (order of magnitude)
- Radiator area: ~2,000–5,000 m² per MW of sustained electrical input

Radiators are segmented, replaceable, stowable for launch, and derived from ASI tanker and station thermal design heritage.

4. Power Generation and Storage

Phase 1 (baseline configuration):

- Deployable solar arrays
- 1–5 MW class electrical input per node
- Lithium-ion or solid-state battery buffering for eclipse periods

Phase 2 (architecture-compatible extension):

- Nuclear heat source with Brayton conversion
- Shared radiator farms
- Power export to adjacent Aegis infrastructure

Power architecture is decoupled from compute payload to allow upgrades without redesign.

5. Radiation Mitigation and Shielding

- Passive aluminum hull shielding
- Targeted water shielding around compute racks
- Optional use of lunar-sourced water in later phases

This approach reduces SEU rates and thermal variability without imposing the mass penalties associated with deep-space hardening.

6. Communications Architecture

Space-to-space:

- Optical (laser) inter-satellite links
- High-bandwidth, low-interference backbone

Space-to-ground:

- Hybrid optical and Ka-band RF
- Optical downlinks to dedicated clear-sky ground stations
- RF fallback for availability and legacy integration

AOCN is optimized for batch, streaming, and asynchronous workloads. The architecture does not assume ultra-low-latency interactive access.

Representative v1.0 Node Specifications

Parameter	Order-of-Magnitude Value
Total Mass	40–80 metric tons
Continuous Compute Output	200–500 kW
Electrical Input	1–2 MW
Radiator Area	2,000–5,000 m^2
Shielding Mass	15–50 tons (water + structure)
Operational Life	10–15 years

Servicing	Robotic + crew-compatible

Radiator area corresponds to approximately **0.37–0.93 American football fields** (5,350 m² per field, including end zones), included here for physical scale reference.

Operational Concept

- Nodes are launched in modular segments
- Assembled robotically or via crewed servicing missions
- Compute payloads are swappable without de-orbit
- Radiators and solar wings are replaceable independently
- Nodes operate autonomously with periodic ground tasking

This operational approach supports incremental capacity growth, rolling hardware refresh, and bounded failure modes without requiring full system replacement.

Representative Workload Classes

The following examples are illustrative rather than exhaustive and characterize workload types compatible with the architectural assumptions above:

- AI and ML training runs
- Climate and physics simulations
- High-energy rendering and modeling
- Secure government compute workloads
- Space-based autonomy and navigation processing
- Pre-processing for Earth-based data centers

AOCN does not compete with terrestrial hyperscalers; it complements them where Earth-side power, cooling, or security constraints dominate.

Integration with the Aegis Ecosystem

AOCN slots naturally into ASI's broader architecture:

- **Tankers:** delivery of water shielding and thermal working fluids
- **Long-Hauler:** logistics, relocation, and servicing
- **Aegis Station:** crewed maintenance and future power sharing
- **Gradient One:** adjacent artificial-gravity research
- **LMM:** space-adjacent AI workloads and autonomy research

In later phases, AOCN functions as a distributed compute layer supporting orbital and cislunar infrastructure.

Architectural Summary

In orbit, compute hardware is not the primary scaling constraint. Power generation and heat rejection dominate system size, mass, and operational complexity.

AOCN explicitly treats these constraints as first-order design drivers and organizes compute capability around infrastructure that can be assembled, serviced, and expanded using current spaceflight technologies.